High Quality Image Formation by Nonlocal Means Applied to High-Angle Annular Darkfield Scanning Transmission Electron Microscopy (HAADF–STEM)*

Peter Binev, Francisco Blanco-Silva, Douglas Blom, Wolfgang Dahmen, Philipp Lamby, Robert Sharpley, and Thomas Vogt

Bericht Nr. 317 **Dezember 2010**

Institut für Geometrie und Praktische Mathematik
RWTH Aachen

Templergraben 55, D–52056 Aachen (Germany)

* This research was supported in part by the College of Arts and Sciences at the University of South Carolina, the Leibniz program of the German Research Foundation, MURI ARO Grant No. W911NF–07–1–0185, and NSF Grant No. DMS–0915104.

High Quality Image Formation by Nonlocal Means Applied to High-Angle Annular Darkfield Scanning Transmission Electron Microscopy (HAADF–STEM) *

Peter Binev, Francisco Blanco-Silva, Douglas Blom, Wolfgang Dahmen, Philipp Lamby, Robert Sharpley, and Thomas Vogt

December 31, 2010

Abstract

We outline a new systematic approach to extracting high quality information from HAADF–STEM images which will be beneficial to the characterization of beam sensitive materials. The idea is to treat *several, possibly many* low electron dose images with specially adapted digital image processing concepts at a minimum allowable spatial resolution. Our goal is to keep the overall cumulative electron dose as low as possible while still staying close to an acceptable level of physical resolution. We shall present the main conceptual imaging concepts and restoration methods that we believe are suitable for carrying out such a program and, in particular, allow one to correct special acquisition artifacts which result in blurring, aliasing, rastering distortions and noise.

1 Introduction

Modern electron microscopic imaging has reached resolutions significantly better than 100 *pm* which allows for unprecedented measurements of the composition and structure of materials [10, 6, 16]. However, one faces several severe obstacles to fully exploiting the information provided by aberration-corrected instruments. On the one hand, one needs to constantly remediate and reduce environmental perturbations such as air flow, acoustic noise, floor vibrations, AC and DC magnetic fields, and temperature fluctuations. On the other hand, high resolution and a good signal to noise ratio requires a high density of electrons per square nanometer. Unfortunately, soft materials are very susceptible to beam damage, and can only be visualized with low dose techniques, resulting in poor resolution and a prohibitively low signal to noise ratio [4]. Our goal is therefore to compensate for the required lower dose by using more sophisticated image processing techniques applied to multiple samples in order

*This research was supported in part by the College of Arts and Sciences at the University of South Carolina, the Leibniz program of the German Research Foundation, MURI ARO Grant # W911NF-07-1-0185, and NSF Grant # DMS-0915104.

to raise the signal to noise ratio necessary for reliable image formation. Preliminary methods and results were reported in [2].

The paper is organized as follows. In Section 2 we begin with briefly describing the standard image formation process in STEM and identify certain factors that affect image quality and resolution. In Section 3 we describe the general problems of using time series of low dosage micrographs in order to reconstruct high quality micrographs. Section 4 has a description of the method of nonlocal means and the variants we use in our algorithms for analysis and processing. Here we apply our methods to a time series of low dose micrographs of the M1 catalyst. In this case beam damage, local jitter and global drifts are relatively small and the expected improvements from our methods are observed. In Section 5 we consider the more challenging case of beam sensitive materials by applying the methods to samples from the class of Zeolites. Finally, in Section 6 we summarize our results and draw some conclusions which will guide our future studies.

2 STEM Imaging

Images produced by electron microscopes offer only an indirect reflection of reality. One measures the distribution of the intensity of electron scattering at a detector. These intensities depend upon the structure and composition of the sample, the information transfer properties of the microscope as well as uncontrolled perturbations by external stimuli. An example of environmental noise due to airflow in the vicinity of the microscope during image acquisition is illustrated in Fig. 1 where the resulting perturbations are reflected in the micrograph. For the image on the left side of the figure, airflow is reducing the contrast and resolution of a dumbbell-pattern obtained by imaging Si along a crystallographic $\langle 110 \rangle$ direction, as well as introducing distortions during the rastering. For the image on the right hand side, the airflow has been turned off, thereby improving the quality of the micrograph. The distortion mainly appears as a spatial and structural change. We are also capable of measuring the sound pressure level in the room as shown in Fig. 1 c) as well as the vibrational and magnetic characteristics. This type of auxiliary information will be useful in developing *similarity checks* in the NLM process described in later sections.

We emphasize that we do *not* attempt to develop techniques that aim at reaching a resolution that is higher than the one permitted by the hardware, but instead aim to recover the level of resolution set by the microscope by only using a *time series* of lower resolution - viz. lower dose - images.

The guiding aspects for our approach can be summarized as follows: Rastering of the beam across the sample enables certain electron imaging and spectroscopic techniques such as mapping by energy dispersive X-ray (EDX) spectroscopy, electron energy loss spectroscopy (EELS) and annular dark-field imaging (ADF). These signals can be obtained simultaneously, allowing direct correlation of image and spectroscopic data. By using a STEM and a high-angle annular detector, it is possible to obtain atomic resolution images where the contrast is directly related to the atomic number ($\approx Z^2$) [5, 13, 8]. This is in contrast to conventional high resolution electron microscopy, which uses phase-contrast, and therefore produces results which need simulation to aid in interpretation. As for beam sensitivity, a critical issue in electron microscopy is the amount of dose needed to produce an image.

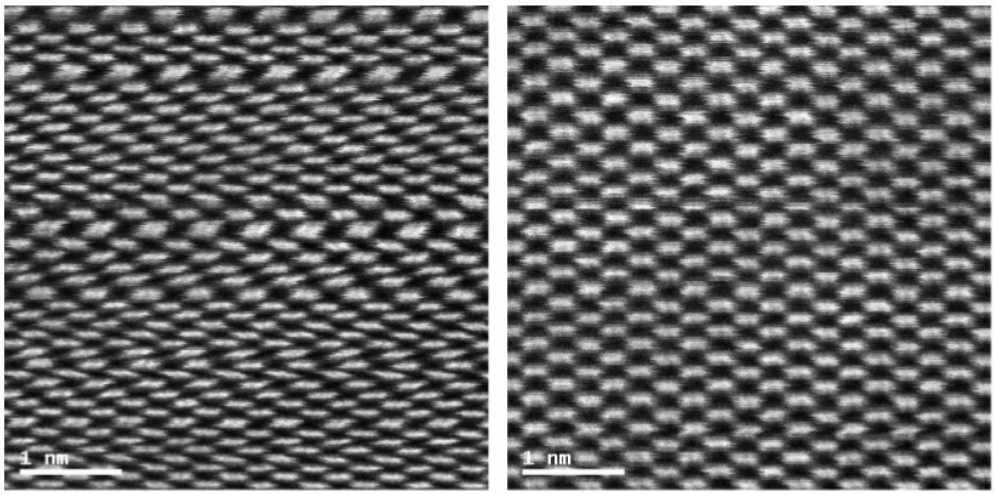

(a) (b)

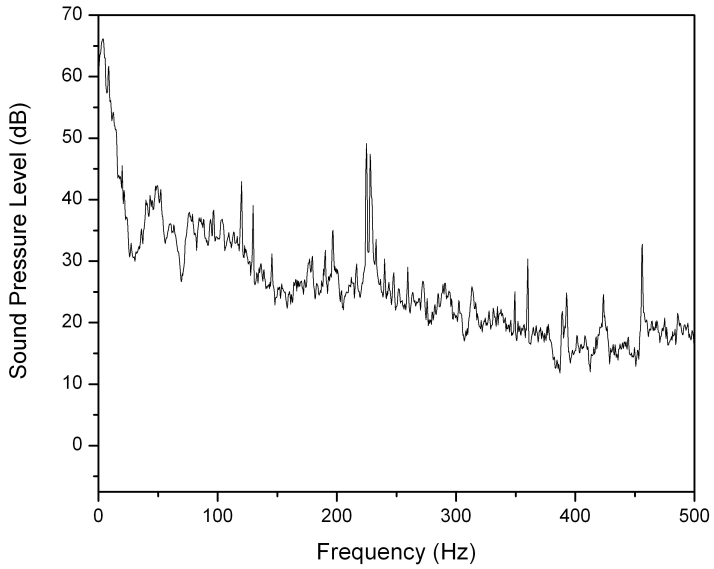

(c)

Figure 1: (a) Si ⟨110⟩ zone axis HAADF STEM micrograph reflecting distortions due to external air pressure perturbations; (b) the airflow is turned off and the location of the Si atomic columns is represented more accurately: (c) sound pressure level (dB) at different frequencies (Hz). Micrographs taken with an exposure of 200 μs per pixel.

Higher dose scans can damage the specimen while lower dose scans result in very low signal to noise ratio. In STEM mode, the electron dose onto the sample can be controlled in a variety of ways. The number of electrons per unit time can be varied by changing the demagnification of the electron source through the strength of the first condenser lens. The dwell time of the probe is typically varied between 7μs and 64μs per pixel in practice, although a much larger range is possible. The size of the image can be varied from a very small number of pixels in a frame (256×256) to over 64 million pixels per image (8192×8192). Finally, the magnification of the image sets the area of the specimen exposed to the electrons and thereby affects the dose per unit area onto the specimen.

3 Formation of High Quality Images from Low Resolution/Noisy Images

Let us briefly recall the standard way of producing high quality images from a series of low resolution/noisy frames. Several observation models that relate the original high-resolution images to the observed low-resolution frames have been proposed in the literature [9]. These are classically formulated as a global model (with local noise $\boldsymbol{n}$) of the form

$$\boldsymbol{y}_t = (D \cdot B_t \cdot M_t)\boldsymbol{x} + \boldsymbol{n}_t, \tag{1}$$

where $\boldsymbol{x}$ is the desired high-resolution image of the sample which is assumed constant during the acquisition of the multiple micrographs, except for any motion and degradation allowed by the model. Therefore, the observed low-resolution images are regarded as the result from warping (M_t), blurring (B_t), and subsampling (D) the original image $\boldsymbol{x}$ and corruption by additive noise $\boldsymbol{n}_t$. Reconstructing the original image $\mathbf{x}$ from observations $\boldsymbol{y}_t$ leads then to an inverse, typically ill-posed problem. However, for STEM imaging this paradigm is hardly applicable because an accurate estimation of the operator M_t is very problematic. The scanning process takes time during which the specimen moves due to electromagnetic, mechanical, or acoustic perturbations. The overall resulting motion may be significant, even for a single frame, but all the more so when taking longer time series of images of the same specimen. Moreover, this motion is very complex. A global drift is typically overlaid by jitter as illustrated by Fig. 2, see also the description in the figure's caption.

Finally, one has to consider the highly non-linear, even non-continuous effects due to the rastering process, which can cause shearing between consecutive rows of pixels in the micrograph. (This is obviously not an issue for standard photography where every pixel value is measured at the same time.) Hence, we conclude, that tracking and estimating the warping by a sufficiently accurate model M_t in (1) is not feasible. A new concept for recovering high quality images from a series of noisy images is therefore required in the case of STEM images.

In the next section we propose an alternative strategy using a variant of *nonlocal* means which needs only an *approximate*, moderately accurate registration and motion tracking, which basically is only needed to estimate the global, large scale drift. Due to the difficulty of the task we see the need to validate our strategy by experiments with materials that exhibit very little beam sensitivity. In particular, inorganic materials allow us to compare

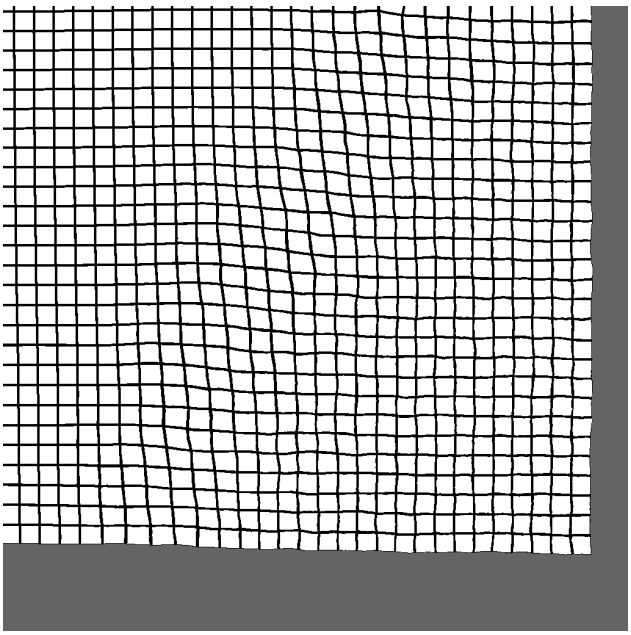

Figure 2: An illustration of the local frame-to-frame distortion. The distortion mapping is first estimated from the global registration of the frames 1 and 9 of the zeolite time series used in Section 5 and then applied to an image of a Cartesian grid to illustrate the complex motion involved. This motion exhibits local jitter overlaid on a global drift upwards and to the left (resulting in the gray region where frame 9 does not overlap the specimen portion depicted by frame 1).

a reconstructed image from low resolution images with a high resolution counterpart of the same object. These experiments are then followed by similar experiments involving more beam-sensitive materials where higher resolution images of these materials are not available due to the resulting beam damage.

Therefore we focus first on inorganic materials which we understand well and that have proven to be stable under HAADF-STEM conditions (see e.g. [15]). In particular, the M1 catalyst, an Mo-V-Te-Nb-oxide, shown in Fig. 3, has various properties that lend themselves to our initial investigations: (1) it has well-understood contrast variations along the $\langle 001 \rangle$ projection, (2) beam-sensitive Te contained in pores of the metal oxide framework can be used to monitor electron beam-induced damage over time series while the surrounding structure does not deteriorate, and (3) defects that can be used as fiducials.

For example, in Fig. 3, a white oval is drawn to show pores in the metal oxide framework containing Te, whose evaporation can be used to monitor long term exposure to electron beams. Thus, measuring time series of M1 at lower resolutions allows us to compare the reconstructed images with micrographs taken at higher resolutions and thereby validate our algorithms and theoretical approaches which guide the treatment of more and more beam sensitive materials.

Of course, one would be able to reduce beam damage (in expectation) if the total accumulated dose used to produce several low resolution images could be kept even below the dose

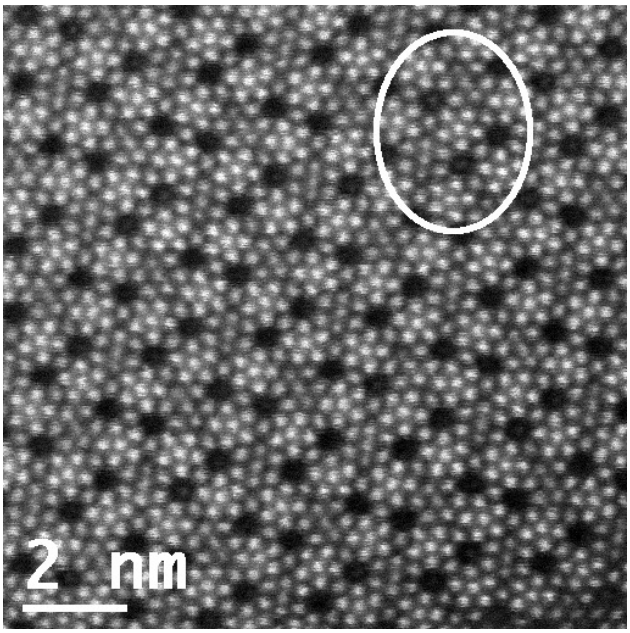

Figure 3: High-resolution HAADF STEM micrograph of the M1 catalyst.

needed for a single high resolution image while still recovering the same information from the low resolution images. But even if in both scenarios the same total dose was necessary, the damage due to heating effects would clearly be smaller when taking successive low dose images. Whether a temporal stretching also has a beneficial relaxation effect on the other sources of beam damage is an open question on which the intended research may actually shed some light. One might note, however, that such principal advantages might come at the price of larger image aquisition times even increasing the movement of the specimen.

4 Nonlocal Means Algorithms for Sequences of Micrographs

Motivated by our earlier observations, we propose an alternative strategy for micrograph image reconstruction based on the non-local means paradigm which has been introduced in [3].

4.1 Nonlocal Means for Time Series

As before, a high quality image is to be recovered from a time series of HAADF STEM micrographs $\boldsymbol{y}_t$ of the same object, where the "time" t is the frame index and runs through a finite set T. Such image assembly algorithms are based on averaging the same specimen portion appearing in different frames. As explained above it is difficult to identify such portions from the noisy low dose frames. It is therefore crucial to employ an averaging technique that is robust with respect to inaccuracies in registration and motion tracking. The concept

of *nonlocal means*, developed by Buades, Coll and Morel in [3] as a denoising algorithm, offers this property. The key point is to assign a higher weight in the averaging process to those patches whose intensity distributions are close to each other and hence more likely to represent the same part of the specimen. Moreover, when the images exhibit repetitive patterns, the denoising effect of averaging can even take advantage of a high similarity of image portions located far apart from each other. The essence of such a procedure can be described as follows.

With every pixel position p in a frame t we associate a (search) neighborhood $N(p,t)$ containing p as well as a patch $R(p,t)$ centered at p. Furthermore for every pixel p in a frame t we make a guess which position p' in frame t' depicts the same specimen portion. We wish to produce an updated (target) value $z(p,t)$ at position p in the frame at t from source values $y(q,t')$ at positions q in the neighborhoods $N(p',t')$ by computing

$$z(p,t) = \frac{\sum_{t\in T_t}\sum_{q\in N(p',t')} w(p,q,t,t')y(q,t')}{\sum_{t\in T_t}\sum_{q\in N(p',t')} w(p,q,t,t')} \tag{2}$$

where T_t denotes a "time neighborhood" of t; that is a collection of timewise neighboring frames that are to be taken into account for the averaging process. Here the *weights* $w(p,q,t,t')$ have the form

$$w(p,q,t,t') := \exp\Big\{-\frac{\operatorname{dist}\big(R(p,t),R(q,t')\big)^2}{\lambda^2}\Big\}, \tag{3}$$

where λ is a data dependent filtering parameter.

The weights serve to quantify the *similarity* between two patches; the more similar two patches are, the more likely it is that the two patches represent the same image portion and consequently we give these pixels higher preference in the averaging process. The similarity is derived from the *distance* $\operatorname{dist}\big(R(p,t),R(q,t')\big)$ between two patches. The distance notion is a crucial parameter of such a scheme. In particular, it allows us to incorporate knowledge about data acquisition and special artifacts and build this into the distance formulation through corresponding transforms applied to the patches. For instance, one could formulate distance notions which are invariant under rotations or other rigid motions of the similarity patches or even filter out the shearing effects which are due to the rastering process, see [11]. We postpone the discussion of this issue and are content for the time being with the perhaps simplest version which views the patch $R(p,t)$ as a vector of intensity values and applies the Euclidean norm to compare two patches of half-size P:

$$\operatorname{dist}\big(R(p,t),R(q,t')\big) := \|R(p,t)-R(q,t')\|_2^2 = \sum_{|r|_\infty\le P}(y(p+r,t)-y(q+r,t'))^2. \tag{4}$$

A few comments on the rationale of such schemes are in order. Obviously, in principle, the weight assigned to a source value $y(q,t')$ is larger as the distance between the corresponding intensities for the respective patches is smaller, regardless of the spatial distance between the respective pixel positions. Thus, in contrast to conventional averaging techniques, closeness in the range is emphasized rather than in the domain, thereby enabling tracking of local jitter (see Fig. 4). The search for similar patches is only limited by the search neighborhood

N_p. For denoising purposes N_p is often chosen as the complete frame, i.e., similar patches are deliberately searched for even in parts of the image that are spatially far away from the pixel to be denoised. In this way self-similarities within the frame, provided by the near-periodic structure of the specimens we are considering, are exploited. On the other hand, averaging over too many patches, none of which exhibit a sufficiently high level of mutual similarity, would cause blurring effects while significantly increasing the computational cost. Hence, for faithful image reconstruction that aims at detecting local artifacts or extra-ordinary features it is necessary to spatially restrict the search neighborhood as much as possible and to compare only patches corresponding to the same specimen portion. This latter aspect, however, can only claim priority once a motion-independent denoising process has sufficiently improved the image quality so that spatial registration becomes feasible.

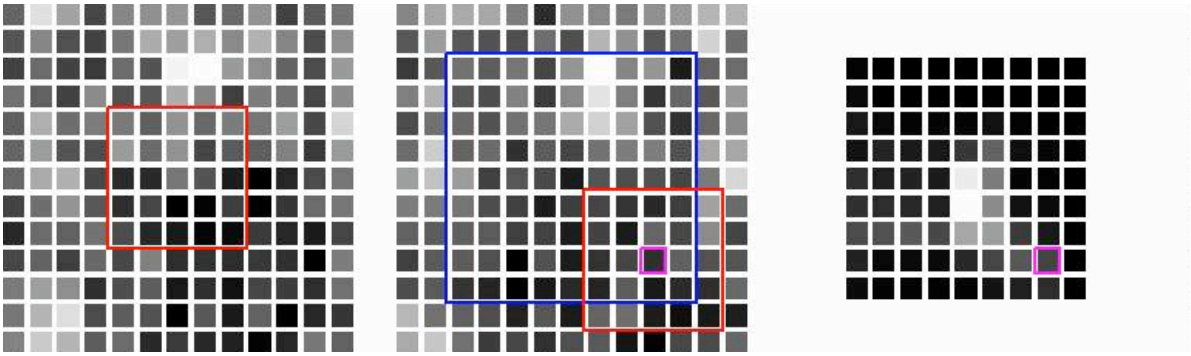

(a) (b) (c)

Figure 4: (a) patch around central pixel (in red); (b) neighborhood (in blue) of central pixel hosting comparison patches; (c) support of weight function for the comparison patches which equals the neighborhood in (b).

While the main issue is to get rid of noise caused by low dose, a limited range of increased spatial resolution can be incorporated in the above framework as well. Concrete algorithms for this task have been developed in [14], but for different types of images. Of course, a concrete scheme based on the above algorithm requires a proper specification of all parameters (patch size/shape, spatial neighborhood size, time neighborhood size, filtering parameter, distance notion). Many of these parameters are found experimentally. Later, during the description of our results, we shall discuss some heuristics.

4.2 A Multi-Stage Algorithm

The preceding discussion already suggests using the nonlocal means averaging process in several stages.

The First Stage: Single Frame Denoising

Recall that the warping that occurs during the image acquisition in HAADF STEM may contain global and local translations, rastering distortion, local rotations, and so on. The overall effect may grow over time and hamper the feature identification in subsequent images. As mentioned earlier a very low signal to noise ratio, increasing distortions or beam

damage in time as well as an unknown complex motion, lower the chance to find sufficiently similar patches in different frames that are timewise far apart. Therefore, at the first stage we employ only a a small time neighborhood T_t (usually consisting only of the frame t itself) and a relatively large spatial neighborhood N_p (usually the whole image) with a simple distance notion such as (4). Actually, this stage is more in the *nonlocal* spirit of the original NLM-algorithm from [3]. The basic idea of this denoising algorithm is to make use of self-similarities within the image itself. As a result one obtains a new time series of smoothed frames in which, however, signals within the micrograph that are not much stronger than the noise level are typically smeared out since the averaging takes too many candidates into account.

The Second Stage: Registration of Denoised Frames

However, the smoothed frames are now better suited for the application of global registration algorithms because the basic structure of the specimen, for instance the positions and shapes of the pores become clearly visible and can reliably be identified. For the experiments in the current work we use the mutual-information-registration code from [7]. This code provides us with maps $(p,t) \to (p',t')$ which are highly accurate so that one can choose very small search neighborhoods in the third stage.

The Third Stage: Multi-Frame Image Formation - Averaging

Now it makes sense to employ more subtle distance notions adapted to the specific features of STEM imaging. Namely, one can replace now the neighborhood $N_p \times N_t$, from which (p',t') is selected, by a (smaller) search domain $\mathcal{N}(p,t)$ that properly takes the frame-to-frame motion into account detected in the first two stages.

Alternative Third Stage: Multi-Frame Image Formation - Median Estimation

An interesting and important alternative to the NLM-type averaging in the multi-frame denoising stage is to determine the target value $z(p,t)$ by computing medians of source pixel values. Median averaging minimizes the distance of the reconstructed image to the source images in the l_1-norm instead of a (weighted) l_2-norm. It has the advantage of being more robust against outliers. Specifically, we set

$$z(p,t) = \text{median}\,\{y(q,t') | t' \in T_t,\ q \in \mathcal{N}(p',t')\} \tag{5}$$

Again it is important to choose an appropriate size of the neighborhoods that are narrowed in space and stretched in time.

In principle the three-stage process can be iterated further with improved similarity criteria. One can gradually decrease the size of spatial neighborhoods while increasing time neighborhoods so as to average eventually only image patches that correspond to each other. It is important to stress though that these iterative passes will always apply to the original data, just using upgraded information concerning the registration extracted from the intermediate frames. In a way, such an iterative procedure may be viewed as gradually refining the image formation in HAADF STEM and modeling the distortions encountered during the imaging process. Moreover, from the possible change of the weights over time one may be able to learn more about beam damage.

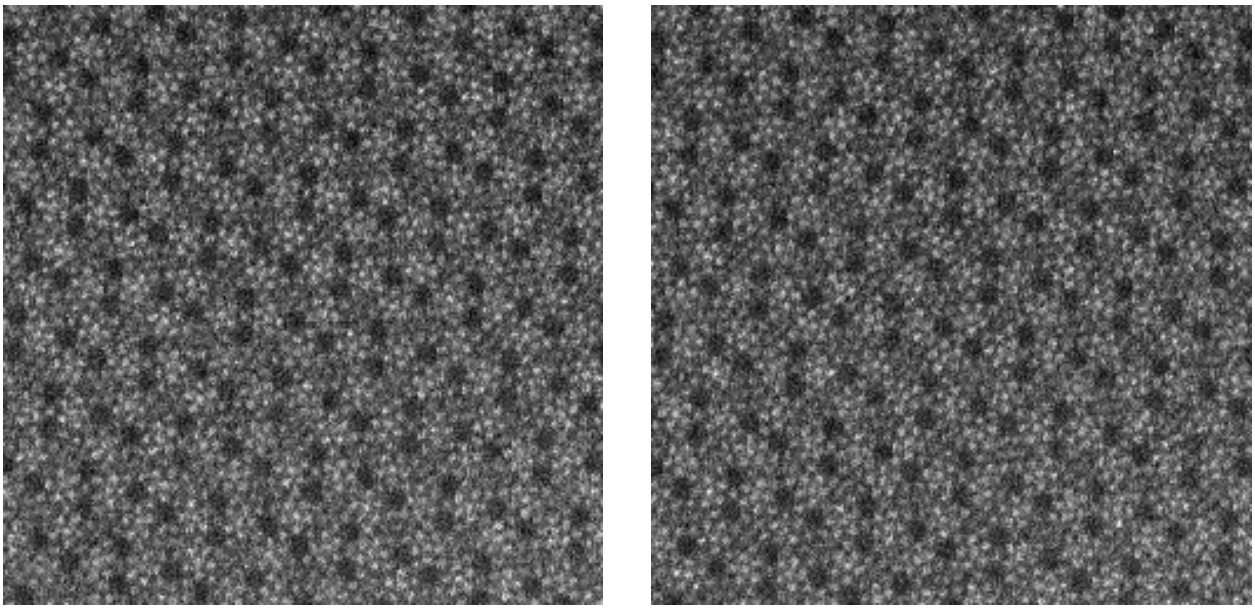

Figure 5: Two samples from a time series of M1-catalyst micrographs.

4.3 M1 Catalyst Micrograph Formation

In the following we apply the program outlined above to a time series of micrographs of the M1 catalyst. The original micrographs have 256×256 pixels, two samples are shown in Fig. 5.

In the first stage we take $\mathcal{N}_p$ as the whole frame. To demonstrate the effect of choosing the parameter λ, we repeat this process twice, both times using a patch size $P = 2$ but with $\lambda = 70,000$, $\lambda = 100,000$ respectively. Choosing between these parameters is done by inspection. A good guess can usually be derived from looking at the difference between the denoised and the noisy image. Assuming that the noise is "white" good parameter settings should give rise to difference images almost without visible structures, see Fig. 6.

A remark concerning the information displayed in the images is in order. The "images" (or better: the data files) contain electron counts registered at the detector after amplification and contain integer values between 0 and about 200,000. In order to display them as images, they are *individually* scaled to the range $[0, 255]$. Intensity changes in the images shown here have their explanation mostly in the fact, that different images might have different maximum values and therefore are scaled differently.

Strictly speaking, the first denoising stage would not have even been mandatory, because the movement of the specimen is generally very small for this particular time series. Within 13 consecutive frames no portion of the specimen moves more than 4 pixels. Therefore we leave the discussion of the registration stage to the next section.

Finally, in Figs. 7 and 8 we form higher quality images using both a similarity driven assembly and a median assembly with a time neighborhood of 11 frames.

In the first case 3×3 pixel spatial neighborhoods were searched, the similarity patches had size 5×5 and the filtering parameter was set to $\lambda = 80,000$. While the result in

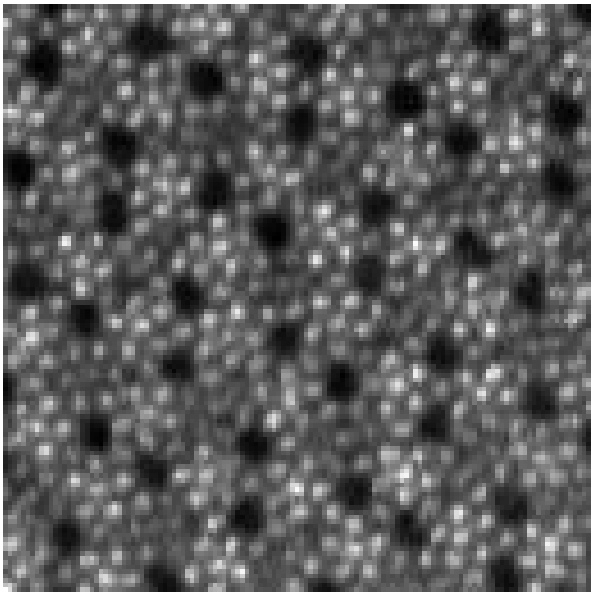

(a) $\lambda = 70,000$

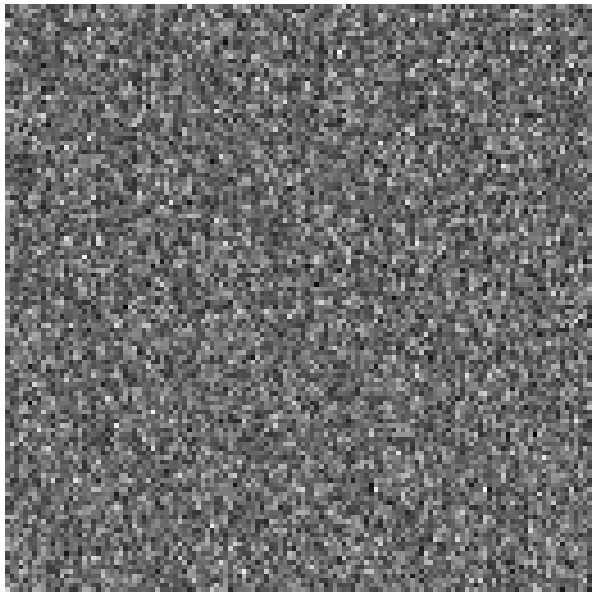

(b) residual for $\lambda = 70,000$

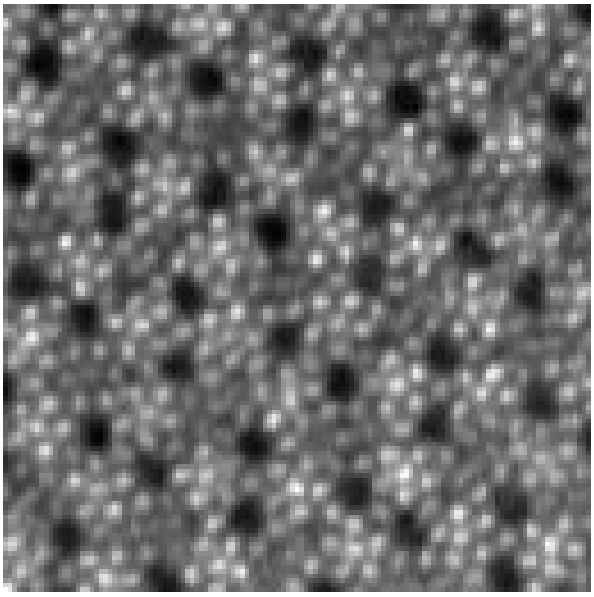

(c) $\lambda = 100,000$

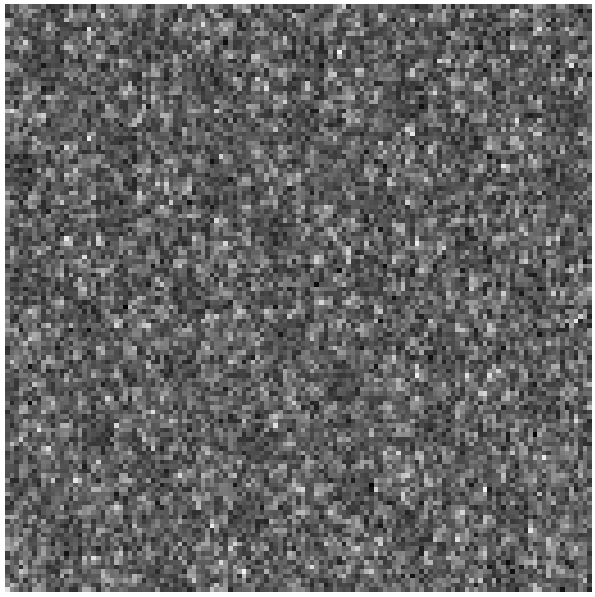

(d) residual for $\lambda = 100,000$

Figure 6: The first frame of the series denoised using the NLM-algorithm with two different sets of parameters ($\lambda = 70,000, 100,000$). The right column shows the differences between the denoised images and the originals. On the bottom row the pores are still clearly visible. Therefore we dismiss this choice of parameters which indeed corresponds to the more blurry denoising result.

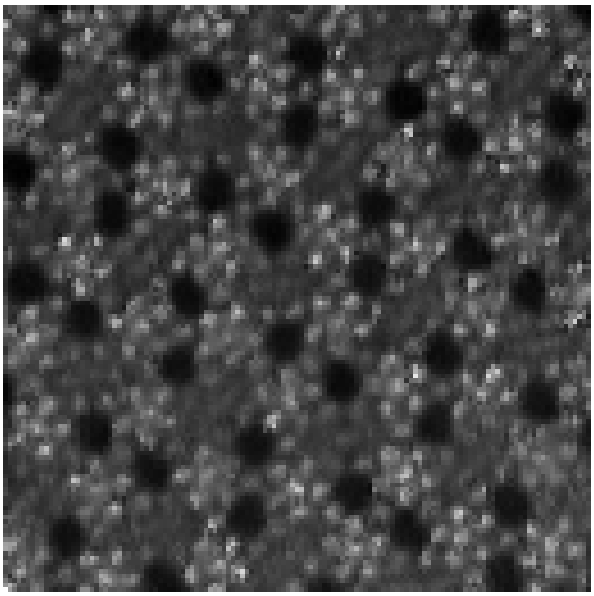

(a) NLM assembly of image

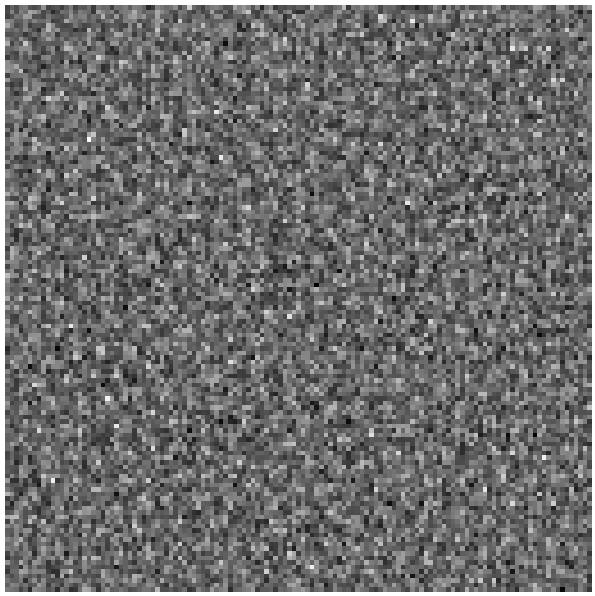

(b) estimated noise with NLM

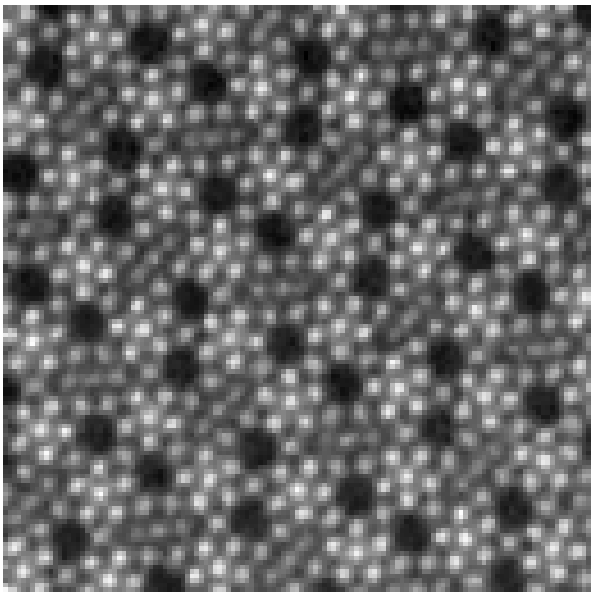

(c) median assembly of image

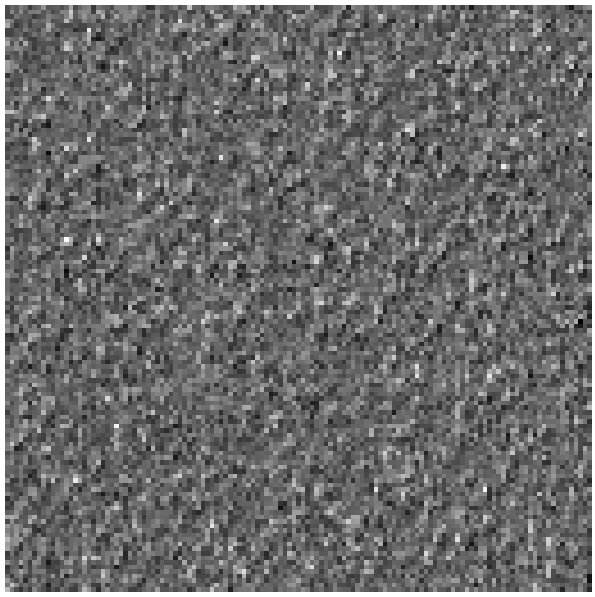

(d) estimated noise with nonlocal median-based approach

Figure 7: Result of assembling 11 images using NLM and a nonlocal median-based approach. The median-based image indicates less averaging out of possibly important information, but its residual in (d) appears to have more unincorporated structure.

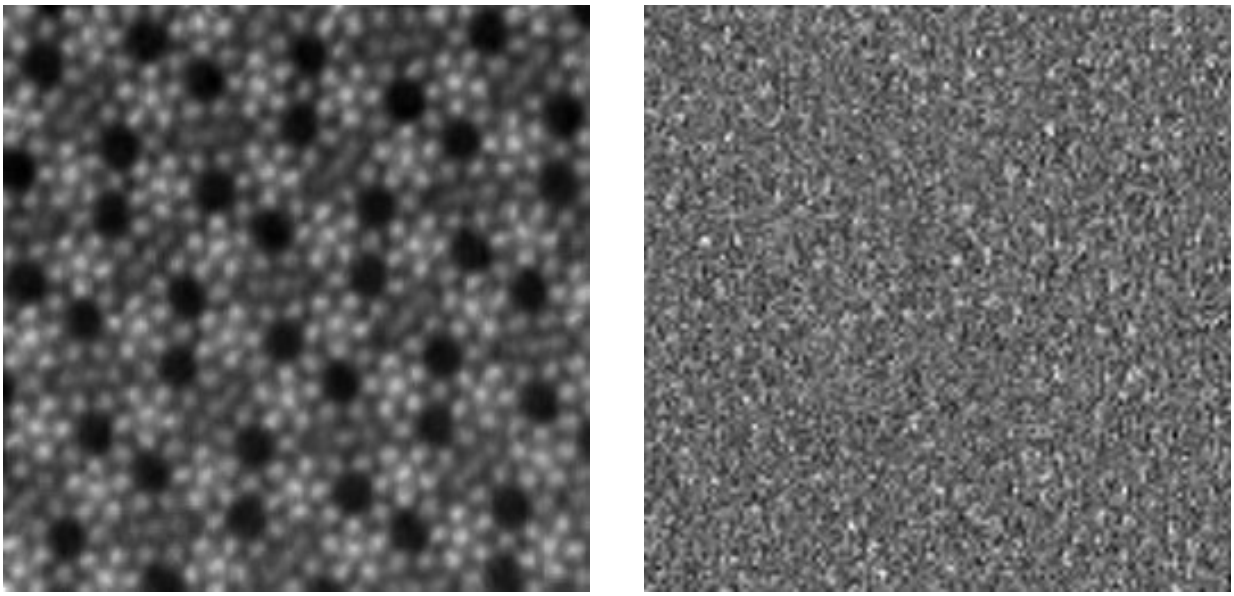

(a) upscaled NLM assembly of image

(b) estimated noise with upscaled NLM

Figure 8: Result of assembling 11 images using upscaled NLM approach. The result is visually better than the one in Fig. 7 (a) but is more blurred than the one in Fig. 7 (c) and eventually misses some detail by smoothing the image too much.

Fig. 7 (a) is received using the standard NLM procedure, we have applied an upscaling technique common in the NLM concept laid out in [14] to receive a better quality image in Fig. 8 (a). However, the upscaling procedure tents to smooth the images which might be an undesirable feature. In the median averaging procedure (shown in Fig. 7 (c)) only 2×2 neighborhoods from each frame were included into the set of pixel values from which to take the median. Images (b) and (d) in the figures show the respective scaled residuals of these methods with frame 1. It should be mentioned that the Fourier transform of the assembled images exhibit the same characteristics as those of the originals.

In general, it seems to us that faint signals, like the ones stemming from Te-atoms contained in the pores (compare with Fig. 3), are more likely be detected by median-assembled images. However, this is subject to further work and validation.

5 Zeolite Micrograph Formation

We conclude this paper with an application of the above strategy to a time series of zeolite micrographs recorded at $2.5 \cdot 10^6$ magnification and taken with a dwell time of 7μs. Zeolites are aluminosilicate materials which contain regular arrays of pores with sizes on the order of many molecular species. They are important materials in a number of absorbtion and catalysis applications. Unfortunately, zeolites are well known to be susceptible to structural collapse under electron beam irradiation. Of key interest for many researchers is the

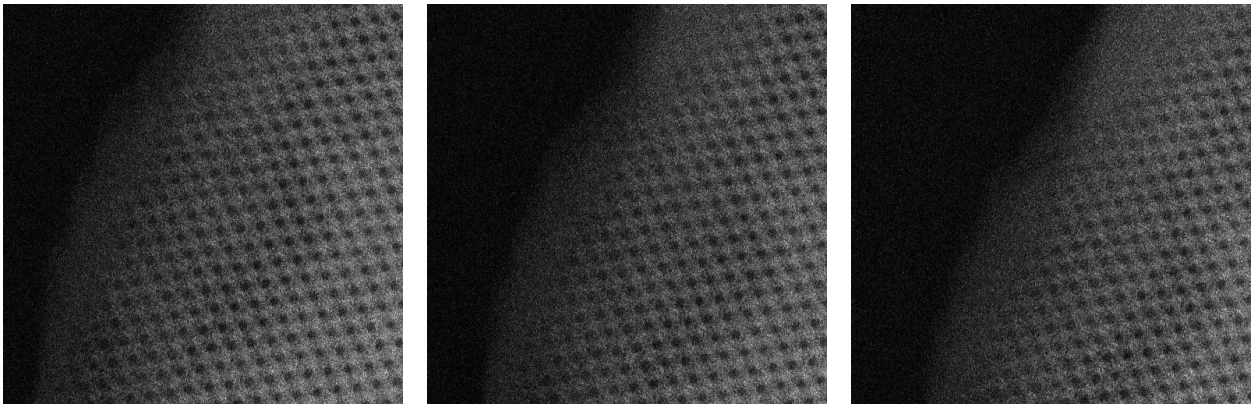

Figure 9: Three original frames: numbers 1, 5, and 8 from the series

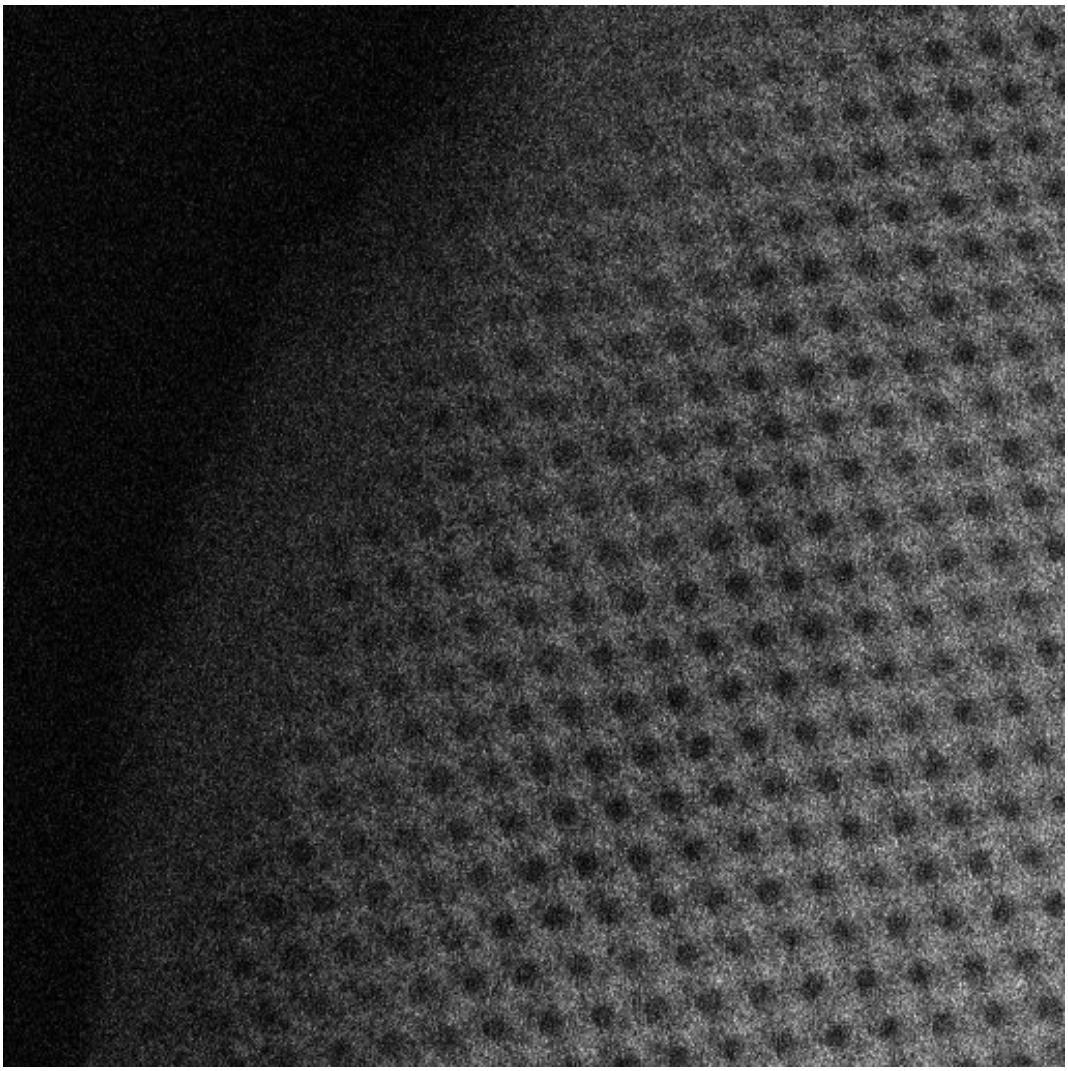

Figure 10: Enlarged rendering of an original zeolite frame 1

arrangement and sizes of the pores in zeolites which are difficult to image via STEM[12].

The original frames have 1024×1024 pixels, but since the multilevel registration code used in stage 2 is more efficient if the pixel width is of the form $2^l + 1$ we cropped the upper left quarter of the images, so that we really work with 513×513 frames. In Fig. 9 we see the first, fifth and eighth frame from this series, and in Fig. 10 an enlarged version of the first frame is shown to present more details. The specimen is wedge-shaped and becomes thicker towards the right side of the image, which expresses itself with increasing intensity values. In the time series one sees that the specimen shifts to the right. Additionally, material is destroyed at the boundary of the wedge.

Stage 1 - In-frame denoising: In this case, denoising before registration is indeed necessary because the originals are too noisy to permit a reliable motion tracking and the deformations occur on a large scale. Fig. 11 shows three denoised frames and Fig. 12 shows an enlarged version of the denoised frame 1.

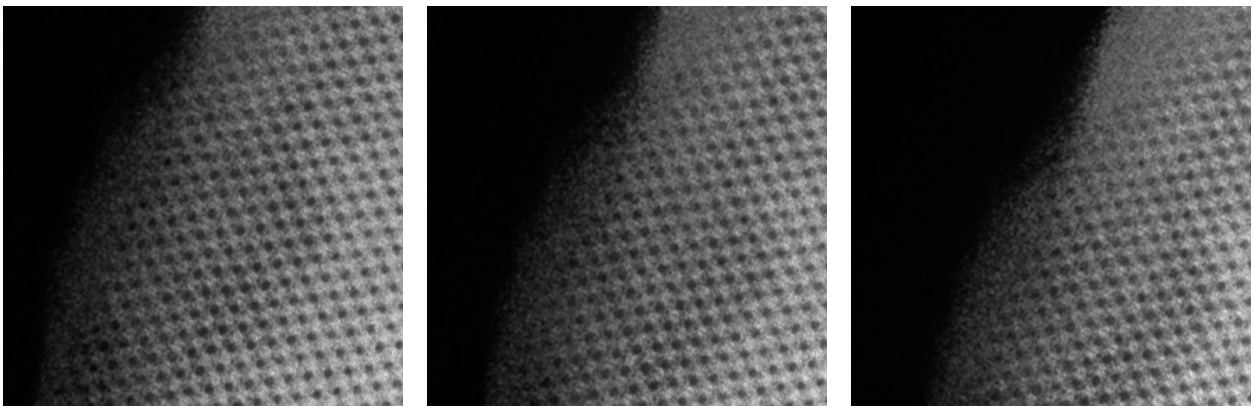

Figure 11: NLM within-frame denoised images for the frames shown in Fig. 9

Stage 2 - Registration: In the second stage we use the denoised frames to register the movement between consecutive frames. For this task we use the mutual-information code by Benjamin Berkels [1, 7]. This code returns for each pixel the information to what position (in fractions of a pixel) in the previous frame it corresponds. The difficulty for the registration is that the rows of the pores look very similar and can easily be confused with each other. The boundary of the specimen is also not a reliable anchor because it degenerates from frame to frame. In Fig. 13 the registration map was used to map consecutive frames onto each other (by some interpolation technique). These maps are also used to validate the correctness of the registration. The images shown here are almost perfectly matched with the frames shown in Fig. 11 if they are superimposed.

Stage 3 - Assembly and Estimation: By composing the maps generated during the registration we can deduce which pixel in the frames 2-9 corresponds to a given pixel in frame 1. We use this information to denoise frame 1, again trying both alternatives (2) and (5).

For the similarity driven assembly we employ 3×3-pixel and for the median assembly 2×2 neighborhood windows. The results are shown in Figs. 14 and 15. Note that in the lower right corner hardly any denoising could be done, because the corresponding pixels have shifted out of the other frames. Here, the median assembly reveals much more details

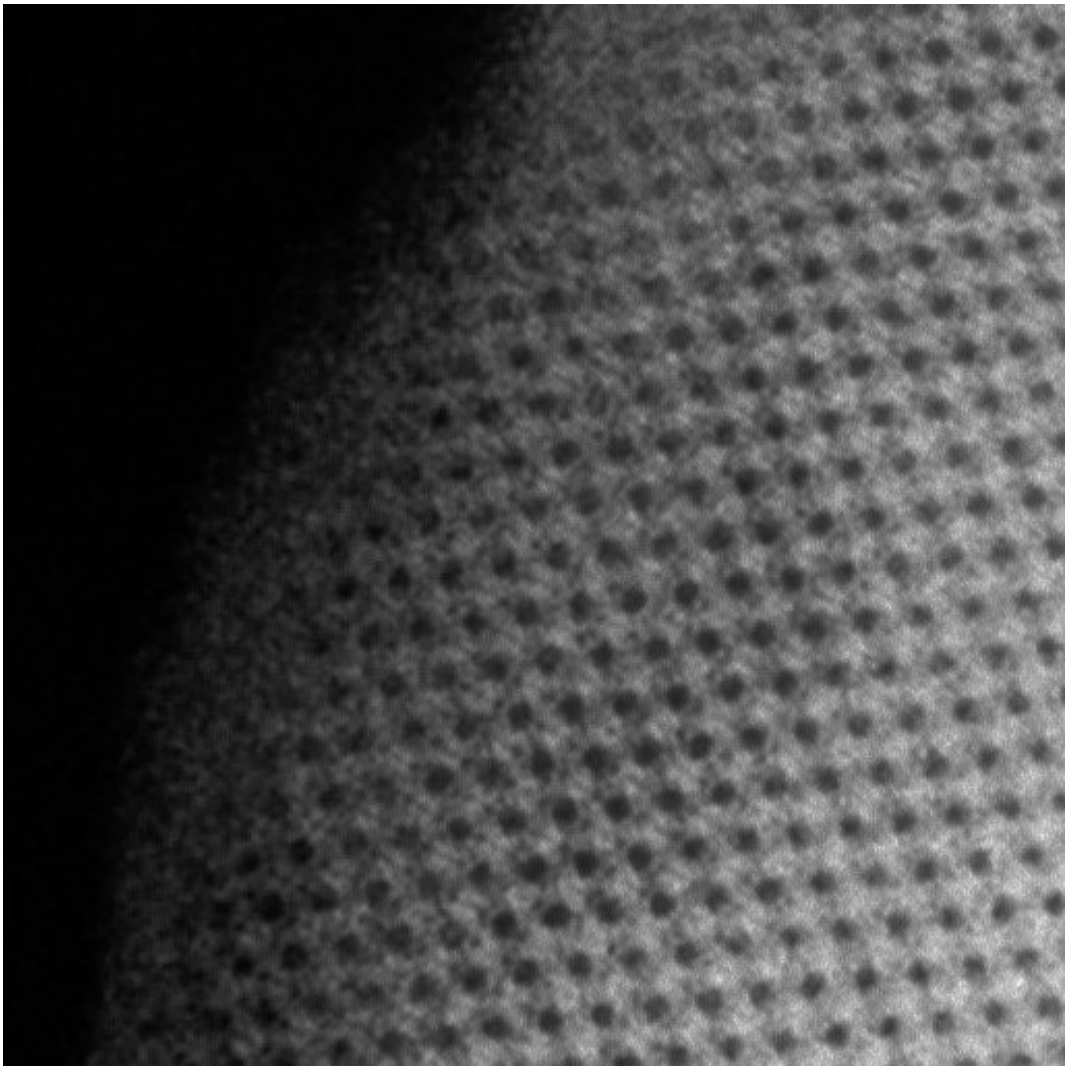

Figure 12: Enlarged rendering of within-frame denoised frame 1

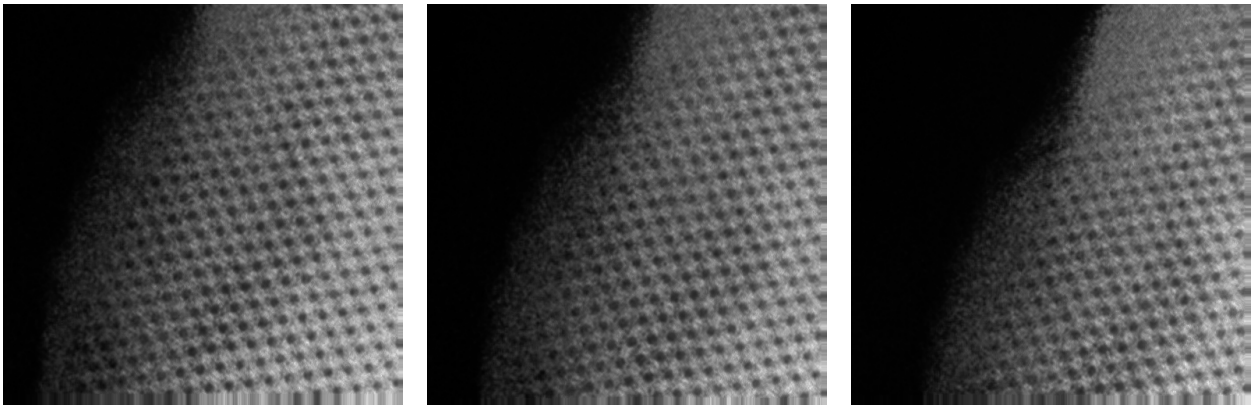

Figure 13: Three examples of the registration: frame 2 mapped onto frame 1, frame 5 mapped onto frame 4 and frame 9 mapped onto frame 8.

than the similarity driven average. The primary pore structure was resolved in even the individual noisy low-dose frames, but following the median assembly, most of the secondary pore structure becomes visible over much of the final assemblage. The in-frame denoised image shows more structure than the NLM time average, probably because it uses much more suitable candidates for averaging due to the ongoing structural collapse of the material upon continued electron irradiation.

Stage 4 - Deblurring: Ideally, it finally remains to deblur the processed images. On one hand, one can try to bring in additional information, for instance, using advanced models for STEM image acquisition. On the other hand, sparse recovery techniques suggest themselves for the corresponding regularization task. Since this concerns work in progress we do not address this issue here any further.

6 Conclusion

We have sketched a new approach to processing STEM images so as to obtain higher quality information from time series of low resolution/low dose frames. Current research focuses on analyzing the effects and identifying suitable choices of the involved scheme parameters. The scheme will then be applied to more and more beam sensitive materials beginning with zeolites. Moreover, we emphasize that the method offers various diagnostic tools. For instance, the variation of the weights over time may shed some light on beam damage mechanisms and their causes. Applying the weights to simple grid test patterns helps to visualize the motion of the specimen during the imaging process for a better understanding.

Acknowledgement. The authors would like to thank Amit Singer and Yoel Shkolnisky for interesting discussions and for introducing them to the method of nonlocal means. We are also indebted to Benjamin Berkels for making his image registration code available to us.

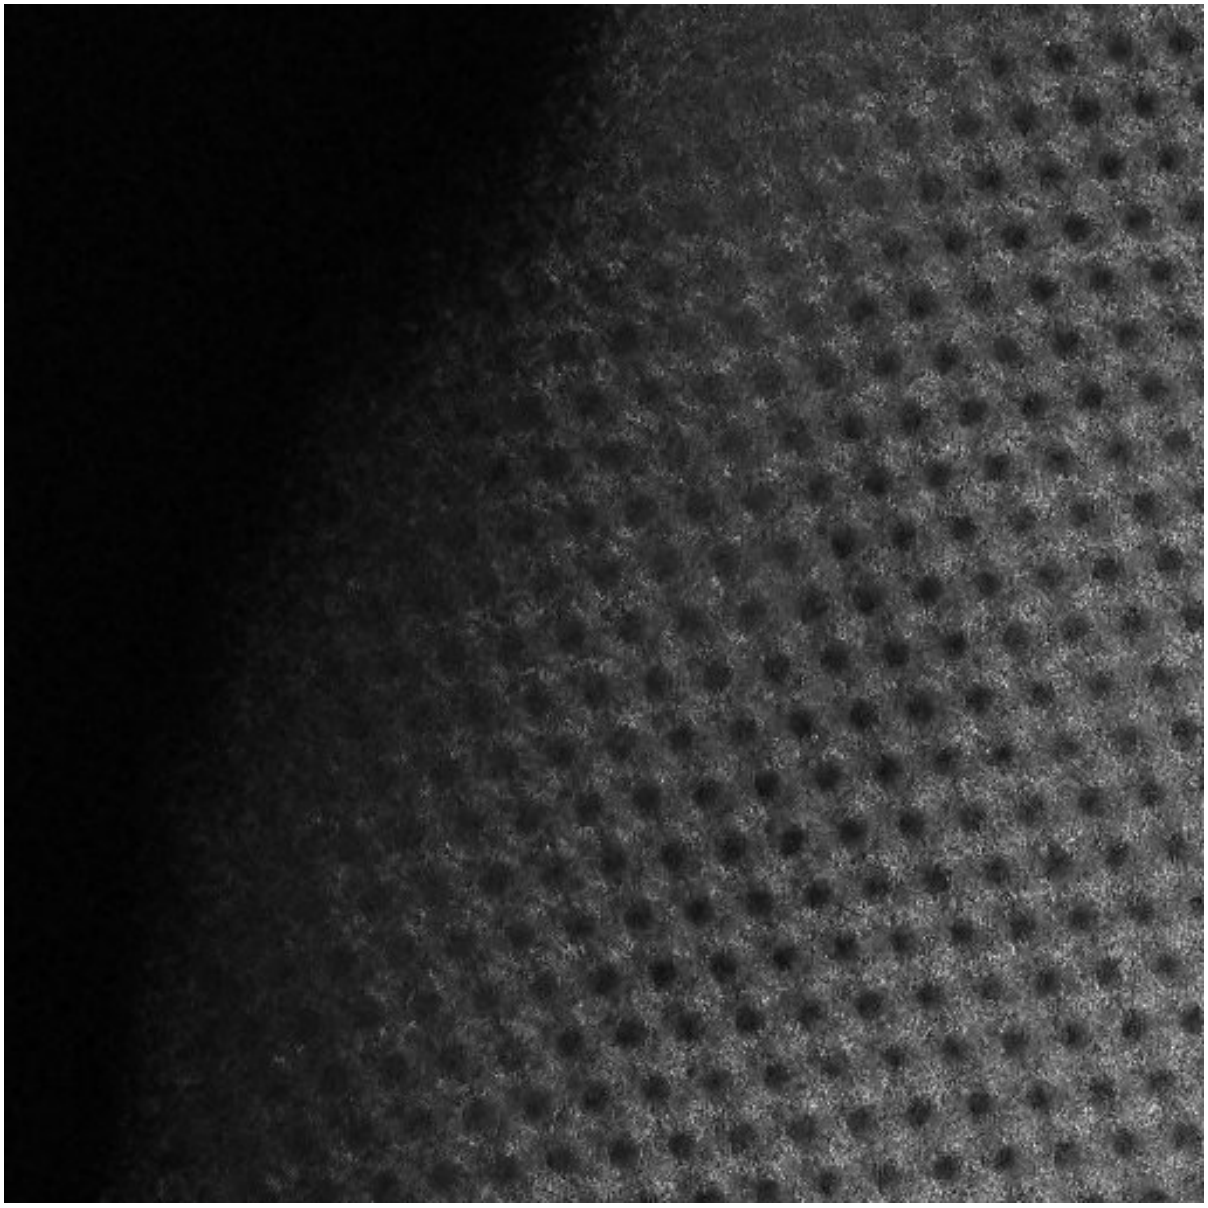

Figure 14: Enlarged rendering of NLM-denoised frame 1, in which the averaging is done only with corresponding registered data in the series of frames.

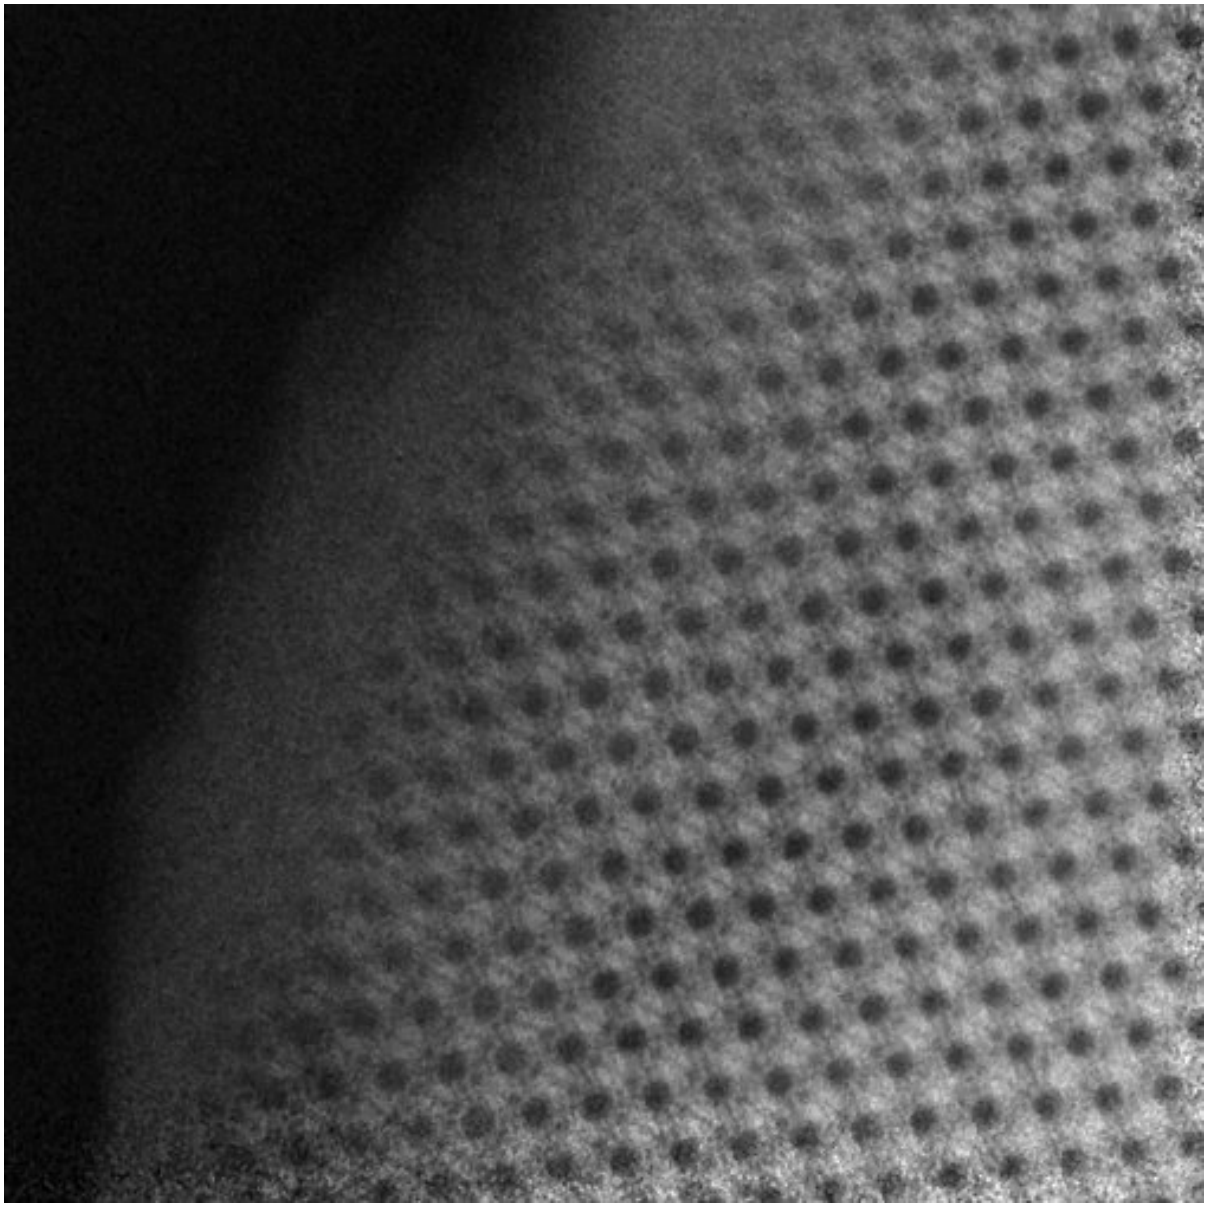

Figure 15: Enlarged rendering of denoised frame 1 using the alternative approach by taking medians of registered frames in the series.

References

[1] Benjamin Berkels, *Joint methods in imaging based on diffuse image representations*, Dissertation, University of Bonn, 2010.

[2] P. Binev, F. Blanco-Silva, D.A. Blom, W. Dahmen, R. Sharpley, and T. Vogt, *Super-Resolution Image Reconstruction by Nonlocal-Means applied to HAADF–STEM*, IMI Preprint Series **2009:06**, University of South Carolina.

[3] A. Buades, B. Coll, and J.M. Morel, *A review of image denoising algorithms with a new one*, Multiscale Modeling and Simulation **4** (2005), 490–530.

[4] J.P. Buban, Q. Ramasse, B. Gipson, N.D. Browning, and H. Stahlberg, *High-resolution low-dose scanning transmission electron microscopy*, J Electron Microsc (Tokyo) **59 (2)** (2010), 103–112.

[5] A.V. Crewe, J. Wall and L. Langmore, *Visibility of single atoms*, Nature **168** (1970), 1338–1340.

[6] R. Erni, M. D. Rossell, C. Kisielowski, and U. Dahmen, *Atomic-Resolution Imaging with a Sub-50-pm Electron Probe*, Physical Review Letters **102** (2009), 096101.1 – 096101.4.

[7] J. Han, B. Berkels, M. Droske, J. Hornegger, M. Rumpf, C. Schaller, J. Scorzin, and H. Urbach, *Mumford-Shah Model for One-to-one Edge Matching*, IEEE Transactions on Image Processing **16:(11)** (2007), 2720–2732.

[8] P. Hartel, H. Rose, and C. Dinges, *Conditions and reasons for incoherent imaging in STEM*, Ultramicroscopy **63** (1996), 93–114.

[9] T.S. Huang and R.Y. Tsai, *Multi-frame image restoration and registration*, Adv. Comput. Vis. Image Process. **1** (1984), 317–339.

[10] C. Kisielowski, B. Freitag, M. Bischoff, H. van Lin, S. Lazar, G. Knippels, P. Tiemeijer, M. van der Stam, S. von Harrach, M. Steckelenburg, et al., *Detection of Single Atoms and Buried Defects in Three Dimensions by Aberration-Corrected Electron Microscope with 0.5-Å Information Limit*, Microscopy and Microanalysis **14** (2008), 469–477.

[11] Markus Navratil, *The New Similarity Notion - Experimental Results*, internal report, 2010.

[12] V. Ortalan, A. Uzun, B.C. Gates, and N.D. Browning, *Direct imaging of single metal atoms and clusters in the pores of dealuminated HY zeolite*, Nature Nanotechnology **5** (2010), 506–510.

[13] S.J. Pennycook and L.A. Boatner, *Chemically Sensitive Structure-image with a Scanning Transmission Electron Microscope*, Science **336** (1988), 565–567.

[14] M. Protter, M. Elad, H. Takeda, and P. Milanfar, *Generalizing the nonlocal-means to super-resolution reconstruction*, IEEE Transactions on Image Processing **18 (1)** (2009), 36–51.

[15] W.D. Pyrz, D.A. Blom, T. Vogt, and D.J. Buttrey, *Direct imaging of the MoVTeNbO M1 phase using an aberration-corrected high-resolution scanning transmission electron microscope*, Angewandte Chemie International Edition **47 (15)** (2008), 2788–2791.

[16] H. Sawada, Y. Tanishiro, N. Ohashi, T. Tomita, F. Hosokawa, T. Kaneyama, Y. Kondo, and K. Takayanagi, *STEM imaging of 47-pm-separated atomic columns by a spherical aberration-corrected electron microscope with a 300-kV cold field emission gun*, J Electron Microsc (Tokyo) **58 (6)** (2009), 357–361.

Peter Binev, Department of Mathematics and the Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208, USA, [binev@math.sc.edu]

Francisco Blanco-Silva, Department of Mathematics and the Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208, USA, [eseprimo@gmail.com]

Douglas A. Blom, Electron Microscopy Center and the NanoCenter, University of South Carolina, Columbia, SC 29208, USA, [Doug.Blom@sc.edu]

Wolfgang Dahmen, Institut für Geometrie und Praktische Mathematik, RWTH Aachen, Templergraben 55, D-52056 Aachen, Germany, [dahmen@igpm.rwth-aachen.de]
and the Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208, USA

Philipp Lamby, Department of Mathematics and the Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208, USA, [lamby@math.sc.edu]

Robert C. Sharpley, Department of Mathematics and the Interdisciplinary Mathematics Institute, University of South Carolina, Columbia, SC 29208, USA, [sharpley@math.sc.edu]

Thomas Vogt, Department of Chemistry and Biochemistry and the NanoCenter, University of South Carolina, Columbia, SC 29208, USA, [TVOGT@mailbox.sc.edu]